Streaming Algorithms for Partitioning Integer Sequences

Christian Konrad[1] and László Kozma[2]

[1] Reykjavik University, Reykjavik, Iceland christiank@ru.is
[2] Universität des Saarlandes, Saarbrücken, Germany kozma@cs.uni-saarland.de

Abstract. We study the problem of partitioning integer sequences in the one-pass data streaming model. Given is an input stream of integers $X \in \{0, 1, \ldots, m\}^n$ of length n with maximum element m, and a parameter p. The goal is to output the positions of separators splitting the input stream into p contiguous blocks such that the maximal weight of a block is minimized. We show that computing an optimal solution requires linear space, and we design space efficient $(1+\epsilon)$-approximation algorithms for this problem following the parametric search framework. We demonstrate that parametric search can be successfully applied in the streaming model, and we present more space efficient refinements of the basic method. All discussed algorithms require space $O(\frac{1}{\epsilon}\,\mathrm{polylog}(m, n, \frac{1}{\epsilon}))$, and we prove that the linear dependency on $\frac{1}{\epsilon}$ is necessary for any possibly randomized one-pass streaming algorithm that computes a $(1+\epsilon)$-approximation.

arXiv:1404.1732v2 [cs.DS] 7 Jul 2014

1 Introduction

In this paper, we study the problem of *partitioning integer sequences*. Given a sequence of integers $X \in \{0, 1, \ldots, m\}^n$ of length n, with maximum element m, and an integer $p \geq 2$, the goal is to partition X into p contiguous blocks such that the maximum weight (sum of the elements) of a block is minimized. In other words, we have to find $p-1$ separators $s_1, \ldots, s_{p-1}$ with $1 = s_0 \leq s_1 \leq \cdots \leq s_{p-1} \leq s_p = n+1$ such that

$$\max \left\{ \sum_{i=s_j}^{s_{j+1}-1} X_i \;\middle|\; j \in \{0, \ldots, p-1\} \right\}$$

is minimized. The value of the previous expression is called the *bottleneck value* of the partitioning. In the following, for any integer $j \in \{0, \ldots, p-1\}$ we refer to the elements $\{X_{s_j}, \ldots, X_{s_{j+1}-1}\}$ as a *partition*, and we refer to the sum of these elements as the *weight* of the partition.

This problem appears in many applications, especially in the context of load balancing, and has been extensively studied both from a theoretical [1–6] and a practical perspective [7, 8]. In the literature, it appears under various names such as *chains-on-chains partitioning* [8, 6] or *1D rectilinear partitioning* [7].

Very efficient exact algorithms for this problem exist, for example the $O(n \log n)$ time algorithm of Khanna et al. [5], the $O(n+p^{1+\epsilon})$ time algorithm of Han et al. [6], and the optimal $O(n)$ time algorithm of Frederickson [9]. However, all existing approaches require either random access to the input or at least multiple access to the same input element. Since in many applications the input integer sequences are huge and cannot be entirely stored in a computer's random access memory, data access is a bottleneck for the previously mentioned algorithms. One example application is the decomposition of computational meshes along space filling curves [10–12]. In parallel scientific computing, for instance in the area of parallel particle simulations or parallel solutions of partial differential equations, huge meshes have to be decomposed and distributed to different computational units. In the space filling curves approach, mesh elements are linearly ordered along a space filling curve which allows the reduction of the multi-dimensional decomposition problem to the one-dimensional problem of partitioning integer sequences, the problem studied in this paper. Today, meshes of Gigabyte or even Terabyte size are common and exceed

by far a computer's random access memory. Algorithms for this problem should therefore have an IO-efficient memory access pattern. In this paper, we are therefore interested in *streaming algorithms* for the problem of partitioning integer sequences.

Streaming Model. In the data streaming model, an algorithm receives its input as a data stream piece by piece. The algorithm is granted a small random access memory which is often only polylogarithmic in the input size. In the present work, we focus on one-pass streaming algorithms, however, depending on the application, an algorithm may be granted multiple passes over the input data in order to further decrease the size of its random access memory. Streaming algorithms find applications in situations where the input data is too large to be stored in local memory and random data access is too costly. For an introduction to streaming algorithms, we refer the reader to [13].

Streaming Algorithms for Partitioning Integer Sequences. We assume that our streaming algorithms receive an input stream $X \in \{0, 1, \ldots, m\}^n$ of length n consisting of integers from the set $\{0, 1, \ldots, m\}$. In addition, we assume that the number of partitions to be created p is stored in the random access memory. All our algorithms make a single pass over the input stream. Since we show that any streaming algorithm that computes an exact solution requires $\Omega(n)$ space, we consider approximation algorithms. We say that an algorithm is a *c-approximation algorithm* if it computes a partitioning with a bottleneck value which is larger than the optimal bottleneck value by at most a factor c. All our algorithms are deterministic. Nevertheless, we prove space lower bounds for possibly *randomized* algorithms. A randomized streaming algorithm is a streaming algorithm that has access to an infinite sequence of random bits, and outputs a correct solution with probability at least $1 - \delta$, for a small constant δ.

We consider the following two variants of the problem:

1. The streaming algorithm outputs separators $s_0, s_1, \ldots, s_p$ that determine the positions of the partitions in the stream. We abbreviate this variant of the problem by Part.
2. The streaming algorithm outputs an upper bound on the bottleneck value of an optimal partitioning. We abbreviate this variant of the problem by PartB (B stands for bottleneck).

There is an important relation between the two variants Part and PartB. A solution to PartB, i.e., a bottleneck value, can be transformed into a solution to Part, i.e., the partition boundaries, via one additional pass over the input stream using the Probe algorithm which is used in many prior works on this problem, e.g. [14, 11, 5]. Probe takes a bottleneck value B and traverses the stream X creating maximal partitions of weight at most B. It is easy to see that Probe succeeds if and only if B is at least as large as the optimal bottleneck value. For this reason, in the definition of PartB, we do not allow a streaming algorithm to output a value that is smaller than the optimal bottleneck value. The Probe algorithm is also an important building block in our work, and we discuss it in more detail in Section 2.

Parametric Search Algorithms. The previously described relation between Part and PartB via the Probe algorithm suggests the application of the *parametric search* framework to this problem, and, in fact, an optimal $O(n)$ time algorithm for this problem is obtained by Frederickson in [9] via this approach. Parametric search was developed by Megiddo more than 30 years ago [15, 16] and has become a standard technique. A parametric search problem is one where the optimal solution is the smallest (or largest) value from a set of candidate solutions of an interval $\{a, a+1, \ldots, b\}$ that passes a certain feasibility test. Usually, monotonicity holds for the values in $\{a, a+1, \ldots, b\}$, i.e., if a value $x \in \{a, a+1, \ldots, b\}$ is feasible then all values $\{x, \ldots, b\}$ (respectively $\{a, \ldots, x\}$) are also feasible. In this situation, using binary search, an $O(\log(b-a)F)$ algorithm can therefore be obtained immediately, where F is the runtime of the feasibility test.

Applied to the problem of partitioning integer sequences, testing feasibility of a value B corresponds to a run of the Probe algorithm. A trivial range for the possible bottleneck values

is $\{1, \ldots, nm\}$ (we discuss better ranges in Section 2), and, therefore, an $O(\log(mn)n)$ time exact algorithm can be obtained. In [9], Frederickson improves this basic idea and obtains an $O(n)$ time algorithm by building data structures on the input sequence that allow the speeding up of the feasibility test, and by exploiting additional information obtained during the feasibility test in order to further narrow down the search space.

Parametric search strongly relies on the fact that the choice of parameter for the next feasibility test depends on the outcome of previous feasibility tests. However, this is impossible to establish in the one-pass streaming model, and, in fact, we prove that in one pass and sublinear space it is impossible to compute the optimal bottleneck value. When relaxing to a $(1+\epsilon)$-approximation, the parametric search framework allows a strategy that results in a one-pass streaming algorithm with space $O(\frac{1}{\epsilon}\log(b-a)S)$, where S is the space required to perform one feasibility test. We run $\Theta(\frac{\log(b-a)}{\epsilon})$ feasibility tests in parallel, testing the values $(1+\epsilon)^i a$ for $i \in \{0, \ldots, \frac{1}{\epsilon}\lceil\log(b-a)\rceil\}$, and we output the smallest parameter of a successful feasibility test. If m, the largest element of the stream, and n, the length of the stream, are known in advance to our algorithm, then in one pass a $(1+\epsilon)$-approximation with space $O(\frac{1}{\epsilon}\log(\frac{mn}{p})p\log(mn))$ can be obtained. Note that this algorithm requires knowledge of the parameters m and n in advance in order to establish a search space that contains the optimal bottleneck value. We regard this algorithm as a baseline algorithm to which we compare our results, and we discuss it in detail in Section 2.

The main contribution of this paper is the design of a new feasibility test: We design the algorithm ProbeExt that takes a parameter B and outputs a feasible value $2^i B$ that is at most by a factor $2+\epsilon$ larger than the optimal bottleneck value, for small ϵ values, if the optimal bottleneck value B^* is at least m/ϵ^2. In some sense, if B^* is sufficiently large, this allows us to run $\Theta(\log(b-a))$ feasibility tests simultaneously. Therefore, compared to the previously described method of running $\Theta(\frac{1}{\epsilon}\log(b-a))$ feasibility tests simultaneously, it is enough to run only $\Theta(\frac{1}{\epsilon})$ of our improved feasibility tests, which improves the space complexity by a $\log(b-a)$ factor. In order to perform our improved feasibility test, we only require knowledge of m in advance while n may be unknown. This is somewhat surprising, since the knowledge of m alone does not allow us to determine an upper limit of the search space for the optimal bottleneck value. Our improved feasibility test, however, can recover from a failed test for x, and continue running a test for some $y > x$, without having to restart the stream. In order to rule out optimal bottleneck values B^* smaller than m/ϵ^2, we additionally run the previously discussed Probe algorithm for bottleneck values in the range $\{m, \ldots, m/\epsilon^2\}$. This allows us to obtain an $O(\frac{1}{\epsilon}\log(\frac{1}{\epsilon})S)$ space algorithm, where S is the space for the ProbeExt algorithm, and the $\log(\frac{1}{\epsilon})$ factor is necessary to rule out cases in which the optimal bottleneck value is smaller than m/ϵ^2.

Which parameters are known in advance? The difficulties of Part and PartB depend strongly on which parameters are known in advance to the algorithm. Suppose that the total weight $S = \sum_i X_i$ of the stream is known in advance. Then it is easy to argue that the optimal bottleneck value B^* is such that $S/p \leq B^* \leq S$. This narrows down the search space, and running $\Theta(\frac{p}{\epsilon})$ copies of the Probe algorithm is enough to obtain a $(1+\epsilon)$-approximation. The knowledge of S provides a lot of information about the input stream. Depending on the application, this may be a reasonable assumption, however, for instance in applications where the weights of elements are estimated on-the-fly, S is certainly not known. Our $(1+\epsilon)$-approximation algorithm that applies our improved parametric search strategy requires only knowledge of m in advance (in fact, any value x with $m \leq x \leq B^*$ will do), while n and S may be unknown. We point out that, in this situation, the initial search space for bottleneck values is unknown since the length of the stream is not known to the algorithm. For the situation where no information

about the parameters is granted in advance, we are only able to obtain a 2-approximation. We leave the existence of a $(1+\epsilon)$-approximation for this situation as an open question.

Communication Complexity. In this paper, we prove two space lower bounds for one-pass streaming algorithms. We show that computing an optimal solution requires $\Omega(n)$ space, and we show that computing a $(1+\epsilon)$-approximation requires $\Omega(\frac{1}{\epsilon}\log n)$ space (for any $\epsilon = \mathrm{O}(n^{1-\gamma})$ for any $\gamma > 0$), showing that the $\frac{1}{\epsilon}$ factor is necessary for obtaining a $(1+\epsilon)$-approximation. Proving space lower bounds for streaming algorithms is often done via communication complexity, and we follow this route in this paper. A one-way two-party communication problem consists of two players, usually denoted by Alice and Bob, who hold inputs Y and Z, respectively. Alice sends a single message to Bob who, upon reception, computes the output of the protocol as a function of Alice's message and his input. The relation to streaming algorithms is as follows: A streaming algorithm for a problem P on data stream $X = Y \circ Z$ (Y concatenated with Z) with space s can be used as a one-way two-party communication protocol for problem P with maximal message size s where player one holds input Y and player two holds input Z. Conversely, a lower bound on the one-way two-party communication complexity of a problem P is also a lower bound on the space requirements for any streaming algorithm for problem P. For an introduction to communication complexity, we refer the reader to [17].

Summary Of Our Results. Our first result is an impossibility result. We show that computing an exact solution to either Part or PartB in one pass requires $\Omega(n)$ space even for randomized algorithms. We therefore study approximation algorithms for the problem. We show that if the maximal value m of the stream X is known in advance, then there is a deterministic $(1+\epsilon)$-approximation algorithm for both Part and PartB using space $\mathrm{O}\left(\frac{1}{\epsilon}\log(\frac{1}{\epsilon})\log(mn^p)\right)$ and $\mathrm{O}\left(\frac{1}{\epsilon}\log(\frac{1}{\epsilon})\log\frac{mn}{\epsilon}\right)$, respectively. These algorithms do not require knowledge of n or of the total weight S of the stream in advance. Then, we consider the hardest case when the algorithm has no information about m, n or S. We design a 2-approximation algorithm for Part using space $\mathrm{O}(p\log(mn))$, and point out a simple 2-approximation algorithm for PartB using space $\mathrm{O}(\log(mn))$. As a counterpoint to these upper bounds, we show that any possibly randomized streaming algorithm that computes a $(1+\epsilon)$-approximation to Part requires $\Omega(\frac{1}{\epsilon}\log n)$ space for any $\epsilon = \mathrm{O}(n^{1-\gamma})$ and any $\gamma > 0$. As our algorithms have a $\frac{1}{\epsilon}\log\frac{1}{\epsilon}$ dependence on ϵ, our lower bound shows that this dependence is optimal up to a logarithmic factor on $\frac{1}{\epsilon}$. Our results are summarized in Figure 1.

m	n	S	Approximation	Space	Remark
Part:					
			exact	$\Omega(n)$	Lower bound (Theorem 6)
-	-	!	$1+\epsilon$	$\mathrm{O}\left(\frac{1}{\epsilon}\log(p)\log(mn^p)\right)$	Baseline (Theorem 1)
!	!	-	$1+\epsilon$	$\mathrm{O}\left(\frac{1}{\epsilon}(p\log^2(n)+\log^2(m))\right)$	Baseline (Theorem 2)
!	-	-	$1+\epsilon$	$\mathrm{O}\left(\frac{1}{\epsilon}\log(\frac{1}{\epsilon})\log(mn^p)\right)$	(Theorem 3)
-	-	-	2	$\mathrm{O}(p\log(mn))$	(Theorem 4)
			$1+\epsilon$	$\Omega(\frac{1}{\epsilon}\log n)$	Lower bound (Theorem 7)
PartB:					
			exact	$\Omega(n)$	Lower bound (Theorem 6)
-	-	!	$1+\epsilon$	$\mathrm{O}\left(\frac{1}{\epsilon}\log(p)\log(mn)\right)$	Baseline (Theorem 1)
!	!	-	$1+\epsilon$	$\mathrm{O}\left(\frac{1}{\epsilon}\log^2(mn)\right)$	Baseline (Theorem 2)
!	-	-	$1+\epsilon$	$\mathrm{O}\left(\frac{1}{\epsilon}\log(\frac{1}{\epsilon})\log(\frac{mn}{\epsilon})\right)$	(Theorem 3)
-	-	-	2	$\mathrm{O}(\log(mn))$	(Theorem 5)

Fig. 1. Overview of our results. In the first three columns we indicate whether advance knowledge of the maximum weight m, the length of the stream n, or the total weight S is required by the algorithm (the ! sign indicates that the respective quantity is required).

Further Related Work. The problem of partitioning integer sequences has been extensively studied in the offline setting, as early as 1988 by Bokhari [1], who presented an exact algorithm with time complexity $O(n^3p)$. Significant progress has since been made on the problem, and the best current algorithm runs in time $O(n)$ independently of p [9]. Previous works use techniques such as dynamic programming, iterative refinement of a partitioning, and parametric search. Most ideas from previous works are not applicable in the streaming model since they require a more flexible data access scheme. The work of Iqbal [14] is closest to our work because it considers approximation algorithms. Furthermore, some of his techniques, such as a parametric search for the optimal bottleneck value, are in their basic features similar to our work. To the best of our knowledge, our work is the first that rigorously follows the parametric search framework in the streaming model.

Outline. First, we discuss the Probe algorithm and we prove the results for our baseline method in Section 2. In Section 3, we discuss the ProbeExt algorithm, which constitutes the main algorithmic technique in this paper. In Section 4, we present algorithms for the case when m is known in advance, and in Section 5, we present algorithms for the case when m is not known in advance. Then, we present our space lower bounds in Section 6. We present our $\Omega(n)$ space lower bound in Subsection 6.1. Then, in Subsection 6.2, we prove a space lower bound for approximation algorithms.

Missing Proofs. Due to space restrictions, many proofs have been moved to the appendix. Lemmas and theorems with deferred proofs are marked with $(*)$.

2 The Probe Algorithm

An important building block for our algorithms is Probe (Algorithm 1), and its variant ProbeB (not explicitly shown). These algorithms have been used in previous works on this problem, e.g. [14]. Probe takes parameters B and p, makes one pass over the input stream and sets up partition separators such that partitions do not exceed a weight of B but are of maximal size. ProbeB performs the same task as Probe, but it does not store the actual separators, and returns only a boolean value indicating whether the algorithm succeeded or failed.

Algorithm 1 Probe(B, p)

```
I ← 1 {current element index}
P ← 1 {current separator index}
W ← 0 {current partition weight}
while input stream not empty do
  if P > p then FAIL end if
  x ← next integer from stream
  I ← I + 1
  if x > B then FAIL end if
  if W + x ≤ B then W ← W + x
  else s_P ← I,  P ← P + 1,  W ← x
  end if
end while
return (1, s_1, ..., s_{p-1}, I)
```

We state now upper and lower bounds on the optimal bottleneck value B^*. Then, we use these bounds in order to derive a bound on the space complexity of Probe and ProbeB. Finally, we show how Probe and ProbeB can be used to obtain a $(1+\epsilon)$-approximation. In the following, let $S = \sum_i X_i$ denote the weight of the entire input integer sequence.

Lemma 1. *Let B^* denote the bottleneck value of an optimal partitioning. Then:*

$$\max\left\{\left\lceil\frac{S}{p}\right\rceil, m\right\} \leq B^* \leq \left\lfloor\frac{S+(p-1)m}{p}\right\rfloor < \left\lfloor\frac{nm}{p}+m\right\rfloor.$$

Proof. For the lower bound $\left\lceil\frac{S}{p}\right\rceil$, observe that the weight of each partition is at most B^*, so their sum is at most $p \cdot B^*$. The integer m is a trivial lower bound since an element of weight m has to be part of some partition.

For the upper bound $\left\lfloor\frac{S+(p-1)m}{p}\right\rfloor$, we construct a partitioning that fulfills this property. Assume that we know S and m in advance. Partition the stream greedily, placing a separator

when the weight of the current partition is at least $B = \frac{S-m}{p}$. The weight of the current partition is thus at most $\lfloor B \rfloor + m = \left\lfloor \frac{S+(p-1)\cdot m}{p} \right\rfloor$. After placing the separators $s_0, \ldots, s_{p-1}$, the sum of the remaining elements (the weight of the last partition) is at most $S - (p-1)B = \frac{S+(p-1)\cdot m}{p}$.

For the upper bound $\left\lfloor \frac{nm}{p} + m \right\rfloor$, note that $S \leq nm$, and therefore, $\left\lfloor \frac{S+(p-1)m}{p} \right\rfloor < \left\lfloor \frac{nm}{p} + m \right\rfloor$. This implies the result. □

The following lemma on the space requirements is easily verifiable and uses the previous bounds on the optimal bottleneck value B^* of Lemma 1.

Lemma 2. PROBE*(B, p) and* PROBEB*(B, p) succeed if and only if the optimal bottleneck value is smaller or equal to B.* PROBE *uses space* $\mathrm{O}(p \log n + \log B + \log m) = \mathrm{O}(\log(mn^p))$ *and* PROBEB *uses space* $\mathrm{O}(\log p + \log B + \log m) = \mathrm{O}(\log(mn))$. □

We show now that if S is known in advance, using Lemma 1, PROBE (respectively PROBEB) can be used to obtain a $(1+\epsilon)$-approximation algorithm for PART (resp. PARTB). As already mentioned in the introduction, this result is obtained by running $\Theta(\frac{\log p}{\log(1+\epsilon)})$ copies of PROBE in parallel. For details, see the proof of Theorem 1 in the appendix.

Theorem 1. *For any positive* $\epsilon = \mathrm{O}(1)$*, if S is known in advance, then by running* $\Theta(\log(p)/\epsilon)$ *copies of* PROBE *(resp.* PROBEB*) we can obtain a* $(1+\epsilon)$*-approximation algorithm for* PART *(resp.* PARTB*). The space requirements are*

- $\mathrm{O}\left(\log(p) \log(mn^p)/\epsilon\right)$ *for* PART*, and*
- $\mathrm{O}\left(\log(p) \log(mn)/\epsilon\right)$ *for* PARTB*.*

Proof. Let $\mathrm{C} = \left\lceil \frac{\log p}{\log(1+\epsilon)} \right\rceil$. We run $C+1$ copies of PROBE (resp. PROBEB) in parallel, with bottleneck values $\frac{S}{p}(1+\epsilon)^i$ for $i \in \{0, 1, \ldots, C\}$. We return the successful partitioning with the smallest bottleneck value. Note that

$$\frac{S}{p}(1+\epsilon)^C \geq \frac{S}{p} \cdot p = S.$$

Let B^* denote the optimal bottleneck value. Since $B^* \leq S$, there is always at least one run of PROBE (resp. PROBEB) that succeeds, due to Lemma 2. Let $B = \frac{S}{p}(1+\epsilon)^{i'}$ be the returned bottleneck value. Suppose first that $i' = 0$. Then $B = S/p$ and since $B^* \geq S/p$, we found the optimum. Otherwise $i' > 0$. Then

$$\frac{S}{p}(1+\epsilon)^{i'-1} \leq B^* \leq \frac{S}{p}(1+\epsilon)^{i'} = B,$$

and therefore $B^*(1+\epsilon) \geq B$ which proves the approximation ratio.

Observe that $C = \Theta(\log(p)/\epsilon)$ for any positive $\epsilon = \mathrm{O}(1)$. The largest bottleneck value for which we run PROBE (resp. PROBEB) is $\mathrm{O}(mn)$. For an upper bound on the total space requirement, we multiply the maximal space requirement of a single copy of PROBE or PROBEB (Lemma 2) by the number of copies $C+1$. The result follows. □

Finally, if S is unknown to the algorithm but m and n are known, then the following holds:

Theorem 2. *For any positive* $\epsilon = \mathrm{O}(1)$*, if m and n are known in advance, then by running* $\Theta(\log(mn)/\epsilon)$ *copies of* PROBE *(resp.* PROBEB*) we can obtain a* $(1+\epsilon)$*-approximation algorithm for* PART *(resp.* PARTB*). The space requirements are* $\mathrm{O}\left((p \log^2 n + \log^2 m)/\epsilon\right)$ *for* PART*, and* $\mathrm{O}\left(\log^2(mn)/\epsilon\right)$ *for* PARTB*.*

The proof of Theorem 2 is omitted since it is essentially equivalent to the proof of Theorem 1 using the initial search space $\{m, m+1, \ldots, mn\}$.

3 The ProbeExt Algorithm

In this section, we present a one-pass streaming algorithm that only requires the knowledge of m in advance. We denote this algorithm by ProbeExt, and similar to the Probe algorithm, we introduce a counterpart ProbeExtB that does not store partition boundaries. ProbeExt receives m and a real number $0 \leq \alpha < 1$ as parameters, and initially tries to set up maximal partitions of size at most $B = m(1+\alpha)$. We discuss the actual purpose of α later, however, we mention that the choice of α does not affect the approximation factor of the algorithm. Let B^* denote the optimal bottleneck value. If $B^* > m(1+\alpha)$ then ProbeExt will reach a state where all p partitions are set up, while there are still integers in the input stream.

Algorithm 2 ProbeExt(m, p, α)

```
I ← 1 {current element index}
P ← 1 {current separator index}
W ← 0 {current parition weight}
B ← m(1 + α) {curr. bottleneck value}
while input stream not empty do
   x ← next integer from stream
   I ← I + 1
   if W + x ≤ B then W ← W + x
   else if P < p then
      s_P ← I,  P ← P + 1,  W ← x
   else {merge adjacent partitions}
      s_P ← I,  B ← 2B,  P ← ⌊p/2⌋ + 1
      for i = 1 ... p/2 do s_i ← s_{2i} end for
      if p even then W ← x
      else W ← W + x end if
   end if
end while
return B, (1, s_1, ..., s_{p-1}, I)
```

In this situation, we *merge* all adjacent partitions i and $i+1$ for odd i. In so doing, we create $\lfloor p/2 \rfloor + 1$ new partitions, each with weight at most $2m(1+\alpha)$. We double the current bottleneck value B from $m(1+\alpha)$ to $2m(1+\alpha)$ and we continue setting up partitions. We perform these steps repeatedly until we reach the end of the stream, and we obtain a bottleneck value of $2^i(1+\alpha)m$, where i denotes the number of merge operations that occurred during the execution of the algorithm. We summarize the space requirements of ProbeExt and ProbeExtB in the following lemma.

Lemma 3. *A run of* ProbeExt *requires space* $\mathrm{O}(\log(mn^p))$*, and a run of* ProbeExtB *requires space* $\mathrm{O}(\log(mn))$*.*

Proof. ProbeExt stores the separators, which accounts for $\mathrm{O}(p \log n)$ space. Furthermore, it stores the variable B which is bounded by the bottleneck value of the partitioning it creates. As we show later that the algorithm is a constant factor approximation, this value is in the order of the optimal bottleneck value, which in turn is bounded by $\mathrm{O}(\frac{mn}{p})$, see Lemma 1. Therefore, we obtain the bound $\mathrm{O}(p \log n + \log(\frac{mn}{p})) = \mathrm{O}(\log(mn^p))$.

ProbeExtB does not store the separators. Therefore, its space requirement is bounded by the optimal bottleneck value $\mathrm{O}(\log(\frac{mn}{p})) = \mathrm{O}(\log(mn))$. □

In the remainder of this section, we show that if the optimal bottleneck value B^* is large compared to m, then the algorithm is close to a 2-approximation (see Lemma 6). We use this fact in Section 4 to obtain a $(1+\epsilon)$-approximation algorithm.

Lemma 4. *Suppose that* ProbeExt *(or* ProbeExtB*) performs i merge operations. Then the weight of the input stream is at least*

$$\frac{pm}{2}\left(2^i(1+\alpha) - \alpha - i\right) - \frac{m}{2}(i+\alpha).$$

Proof. We develop a lower bound on the total weight of the stream after i merge operations have been executed, and we denote this lower bound by $LB(i)$.

Consider the situation just before the first merge operation. Denote by w_j the weight of the jth partition. Note that for all j we have $w_j + w_{j+1} > m(1+\alpha)$, otherwise the algorithm would have created a single partition instead of the two adjacent partitions j and $j+1$. Thus, if p is

even, we have $\sum_j w_j \geq \frac{1}{2}pm(1+\alpha)$ and if p is odd, we have $\sum_j w_j \geq \frac{1}{2}(p-1)m(1+\alpha)$. To unify the analysis for the even and the odd case, we set

$$LB(1) = \frac{1}{2}(p-1)m(1+\alpha).$$

Consider now the situation just before the ith merge operation, again denoting by w_j the weight of the jth partition. If p is even, then the weight of the first $p/2$ partitions is at least $LB(i-1)$. Clearly, each of the remaining $p/2$ partitions have weight of at least $\lceil 2^{i-1}m(1+\alpha) - m \rceil$, hence:

$$\sum_j w_j = \sum_{j=1}^{p/2} w_j + \sum_{j=p/2+1}^{p} w_j \geq LB(i-1) + \frac{p}{2}(2^{i-1}m(1+\alpha) - m).$$

Suppose now that p is odd. Then the weight of the first $(p-1)/2$ partitions is at least $LB(i-1) - \lfloor 2^{i-2}m(1+\alpha) \rfloor$. The remaining $(p+1)/2$ partitions have a weight of at least $\lceil 2^{i-1}m(1+\alpha) - m \rceil$, and we obtain

$$\begin{aligned}\sum_j w_j = \sum_{j=1}^{(p-1)/2} w_j + \sum_{j=(p+1)/2}^{p} w_j &\geq LB(i-1) - \lfloor 2^{i-2}m(1+\alpha) \rfloor + \frac{p+1}{2} \lceil 2^{i-1}m(1+\alpha) - m \rceil \\ &\geq LB(i-1) + \frac{p}{2}(2^{i-1}m(1+\alpha) - m) - \frac{1}{2}m.\end{aligned}$$

In order to treat the even and the odd case at the same time, we set

$$LB(i) = LB(i-1) + \frac{p}{2}(2^{i-1}m(1+\alpha) - m) - \frac{1}{2}m,$$

and we eliminate the recursion:

$$LB(i) = \sum_{j=2}^{i} \big(LB(j) - LB(j-1)\big) + LB(1) = \frac{pm}{2}\left(2^i(1+\alpha) - \alpha - i\right) - \frac{m}{2}(i+\alpha).$$

□

Lemma 5. *Suppose that* ProbeExt *(or* ProbeExtB*) performs i merge operations, for $i \geq 2$. Then for any $0 \leq \alpha < 1$,* ProbeExt *(resp.* ProbeExtB*) has an approximation factor of at most*

$$2 + \frac{2(\alpha+i)}{2^{i-1}(1+\alpha) - i - \alpha}.$$

Proof. Let us denote the bottleneck value of the solution returned by ProbeExt (resp. ProbeExtB) by $B = 2^i m(1+\alpha)$, and let B^* denote the optimal bottleneck value. By Lemma 4, the total weight S of the stream is at least $\frac{pm}{2}\left(2^i(1+\alpha) - \alpha - i\right) - \frac{m}{2}(i+\alpha)$, and B^* is at least a p-fraction of S (Lemma 1). The approximation factor of ProbeExt (resp. ProbeExtB) can be bounded as follows:

$$\frac{B}{B^*} \leq \frac{2^i m(1+\alpha)}{\frac{1}{p} \cdot \left(\frac{pm}{2}\left(2^i(1+\alpha) - \alpha - i\right) - \frac{m}{2}(i+\alpha)\right)} \leq 2 + \frac{2(\alpha+i)}{2^{i-1}(1+\alpha) - i - \alpha}\ .$$

□

We conclude with the following result:

Lemma 6. *For any* $0 \leq \alpha < 1$ PROBEEXT *(or* PROBEEXTB*) is a* $(2+\epsilon)$*-approximation algorithm if the optimal bottleneck value satisfies* $B^* > m/\epsilon^2$*, assuming* $0 < \epsilon \leq 1/64$.

Proof. By Lemma 5, PROBEEXT is a $\left(2 + \frac{2(\alpha+i)}{2^{i-1}(1+\alpha)-i-\alpha}\right)$- approximation algorithm if i merge operations have been executed. We have:

$$\frac{2(\alpha+i)}{2^{i-1}(1+\alpha)-i-\alpha} \leq \frac{2(i+1)}{2^{i-1}-(i+1)} \leq \frac{1}{2^{i/2}} = \epsilon.$$

The first inequality uses the bounds on α. To make sure that all quantities are positive and the second inequality also holds, we require $i \geq 12$. The last equality gives $i = 2\log(\frac{1}{\epsilon})$, and our previous bound on i forces $\epsilon \leq 1/64$. Under these conditions, PROBEEXT is a $(2+\epsilon)$-approximation algorithm. Since $B \geq B^*$, PROBEEXT must perform i merge operations if $B^* > m(1+\alpha)2^{i-1}$. Since $m(1+\alpha)2^{i-1} < m2^i = \frac{m}{\epsilon^2}$, the condition $B^* > m/\epsilon^2$ is a sufficient one. □

4 $(1+\epsilon)$-approximation for Known m

In this section, we present a $(1+\epsilon)$-approximation algorithm for PART and PARTB, using as building blocks the PROBE and PROBEEXT algorithms presented in Sections 2 and 3. We assume that the maximum m of the sequence is known in advance.

Algorithm 3 $(1+\epsilon)$-Approximation For Known m

$\delta \leftarrow \frac{\epsilon}{1+\frac{1}{2}\epsilon}$
do in parallel {in one pass}
  PROBE$(2^i(1+\epsilon)^j m, p)$ for all $i \in \{0,1,\ldots,\lceil\log(1/\delta^2)\rceil\}$, $j \in \left\{0,1,\ldots,\left\lceil\frac{1}{\log(1+\epsilon)}\right\rceil\right\}$
  PROBEEXT$(m, p, (1+\frac{1}{2}\epsilon)^j - 1)$, for all $j \in \left\{0,1,\ldots,\left\lceil\frac{1}{\log(1+\frac{1}{2}\epsilon)}\right\rceil\right\}$
end do
return partitioning with smallest bottleneck value

Algorithm 3 runs multiple copies of the the PROBE algorithm and multiple copies of the PROBEEXT algorithm in parallel. We argue that if the optimal bottleneck value B^* is sufficiently large then one run of the PROBEEXT algorithm will return a $(1+\epsilon)$-approximation. If B^* is small, then a run of the PROBE algorithm will return a $(1+\epsilon)$-approximation.

Theorem 3. *For any* $\epsilon < 1/64$*, Algorithm 3 is a* $(1+\epsilon)$*-approximation streaming algorithm for* PART *using space* $\mathrm{O}\left(\frac{1}{\epsilon}\log(\frac{1}{\epsilon})\log(mn^p)\right)$*. The analogous algorithm for* PARTB *is a* $(1+\epsilon)$*-approximation streaming algorithm using space* $\mathrm{O}\left(\frac{1}{\epsilon}\log(\frac{1}{\epsilon})\log(mn)\right)$.

Proof. We distinguish two cases depending on the magnitude of the optimal bottleneck value B^*. In the following, $\delta = \frac{\epsilon}{1+\frac{1}{2}\epsilon}$ as in Algorithm 3.

1. $B^* \leq m/\delta^2$: We show that one of the runs of PROBE is successful and returns a partitioning with bottleneck value B such that $B \leq B^*(1+\epsilon)$. We run PROBE with bottleneck values $2^i(1+\epsilon)^j m$, and since there is a run with $i = \lceil\log(1/\delta^2)\rceil$, there is at least one successful run of PROBE with a bottleneck value of at most m/δ^2. Let $B = 2^{i'}(1+\epsilon)^{j'}m$ denote the smallest bottleneck value of a successful run for values i', j'. Suppose that $j' > 0$. Then the run with bottleneck value $2^{i'}(1+\epsilon)^{j'-1}m$ failed, and therefore

$$B = 2^{i'}(1+\epsilon)^{j'}m \geq B^* > 2^{i'}(1+\epsilon)^{j'-1}m,$$

which implies $B \leq (1+\epsilon)B^*$. Suppose now that $j' = 0$ and $i' > 0$. Then $B = 2^{i'}m$, and the run with bottleneck value $2^{i'-1}(1+\epsilon)^{\left\lceil \frac{1}{\log(1+\epsilon)} \right\rceil - 1}m$ failed, and therefore

$$B = 2^{i'}m \geq B^* > 2^{i'-1}(1+\epsilon)^{\left\lceil \frac{1}{\log(1+\epsilon)} \right\rceil - 1}m,$$

which also implies $B \leq (1+\epsilon)B^*$. If $i' = j' = 0$, then the algorithm found an optimal solution with bottleneck value m.
Since for $\epsilon = \mathrm{O}(1)$ we have $\log(1+\epsilon) = \mathrm{O}(\epsilon)$ and $\delta = \Theta(\epsilon)$, the space requirement for the runs of PROBE is $\mathrm{O}\left(\frac{1}{\epsilon}\log(\frac{1}{\epsilon})\log(mn^p)\right)$, and if we run PROBEB the space requirement is $\mathrm{O}\left(\frac{1}{\epsilon}\log(\frac{1}{\epsilon})\log(\frac{mp}{\epsilon})\right)$.

2. $B^* > m/\delta^2$: By Lemma 6, PROBEEXT and PROBEEXTB are $(2+\delta)$-approximation algorithms for any α. Let $B = 2^{i'}(1+\frac{1}{2}\epsilon)^{j'}m$ be the smallest value output by any of the PROBEEXT runs, for some values of i' and j'. Suppose that $j' > 0$. Then the run with $j = j' - 1$ reports the bottleneck value $2^{i'+1}(1+\frac{1}{2}\epsilon)^{j'-1}m$. Clearly, it cannot return $2^k(1+\frac{1}{2}\epsilon)^{j'-1}m$ for $k \leq i'$ since B is the smallest returned value. On the other hand, it cannot return a bottleneck with $k \geq i'+2$ since then it would have an approximation ratio larger than $2+\delta$, contradicting Lemma 6.

$$2^{i'+2}(1+\frac{1}{2}\epsilon)^{j'-1}m = \frac{4}{1+\epsilon}\cdot B \geq \frac{4}{1+\epsilon}B^* > (2+\delta)B^*.$$

Thus, the run with $j = j' - 1$ returns the bottleneck value $2^{i'+1}(1+\frac{1}{2}\epsilon)^{j'-1}m$. Since this is a $2+\delta$ approximation, we obtain

$$(2+\delta)B^* \geq 2^{i'+1}(1+\frac{1}{2}\epsilon)^{j'-1}m = \frac{2}{1+\frac{1}{2}\epsilon}B \Rightarrow$$
$$B \leq \frac{(2+\delta)(1+\frac{1}{2}\epsilon)}{2}B^* = (1+\epsilon)B^*.$$

Suppose now that $j' = 0$ and $i' > 0$. Consider the run for $j = \left\lceil \frac{1}{\log(1+\frac{1}{2}\epsilon)} \right\rceil - 1$. By a similar argument as before, the run outputs the bottleneck value $2^{i'}(1+\frac{1}{2}\epsilon)^{\left\lceil \frac{1}{\log(1+\frac{1}{2}\epsilon)} \right\rceil - 1}m = B(1+\frac{1}{2}\epsilon)^{\left\lceil \frac{1}{\log(1+\frac{1}{2}\epsilon)} \right\rceil - 1}$. This implies that

$$B^* \geq \frac{B(1+\frac{1}{2}\epsilon)^{\left\lceil \frac{1}{\log(1+\frac{1}{2}\epsilon)} \right\rceil - 1}}{2+\delta} \geq \frac{B(1+\frac{1}{2}\epsilon)^{\frac{1}{\log(1+\frac{1}{2}\epsilon)} - 1}}{2+\delta},$$

which also implies that $B \leq (1+\epsilon)B^*$. Finally, if $i' = j' = 0$, then the algorithm did not perform a merge operation and found an optimal solution with bottleneck value m.
Since $\frac{1}{\log(1+\frac{1}{2}\epsilon)} = \mathrm{O}(\frac{1}{\epsilon})$, the space requirement for the runs of PROBEEXT is $\mathrm{O}\left(\frac{1}{\epsilon}\log(mn^p)\right)$, and if we run PROBEEXTB the space requirement is $\mathrm{O}\left(\frac{1}{\epsilon}\log(mn)\right)$.

For PART, the space requirements are dominated by the runs of the PROBE algorithm. For PARTB, we obtain space $\mathrm{O}\left(\frac{1}{\epsilon}\log(\frac{1}{\epsilon})\log(\frac{mp}{\epsilon}) + \frac{1}{\epsilon}\log(mn)\right)$, and using $p \leq n$ this simplifies to $\mathrm{O}\left(\frac{1}{\epsilon}\log(\frac{1}{\epsilon})\log(\frac{mn}{\epsilon})\right)$. □

5 Algorithms for Unknown m

In this section, we present simple 2-approximation algorithms for PART (resp. PARTB) that do not require the knowledge of any parameter in advance.

Our algorithm for PART works as follows: Suppose that the algorithm has seen the elements $X_1, \ldots, X_i$ and it has partitioned them into p parts with weights $w_1, \ldots, w_p$. If the algorithm now reads the input $x = X_{i+1}$, it will run the PROBE algorithm on the sequence $w_1, w_2, \ldots, w_p, x$ with a bottleneck value B that is at most twice the optimum for a partitioning of $X_1, \ldots, X_{i+1}$ into p parts. See Algorithm 4 and Theorem 4 for further details. The algorithm for PARTB is even simpler, and is described in Theorem 5.

Algorithm 4 2-Approximation for PART

```
w_i ← 0 for all 1 ≤ i ≤ p {partition weights}
I ← 1 {current element index}
S ← 0 {current total weight of input stream}
m ← 0 {current maximum}
while input stream not empty do
    x ← next integer from stream
    I ← I + 1, S ← S + x, m ← max{m, x}
    B ← 2·max{m, S/p} {update bottleneck value}
    Run PROBE(B) on sequence w_1, w_2, ..., w_p, x,
        and store the new separators s_0, ..., s_p
        and the new partition weights w_1, ..., w_p
end while
return B, (1, s_1, ..., s_{p-1}, I)
```

Theorem 4. *Algorithm 4 is a 2-approximation algorithm for* PART *and uses space* $O(p \log(mn))$.

Proof. First, suppose that the run of PROBE in Algorithm 4 succeeds in every iteration. Then, the last bottleneck value is $B = 2 \cdot \max\{m, S/p\}$. By Lemma 1, we have $\max\{m, S/p\} \leq B^* \leq \max\{m, S/p\} + m$, and since $m \leq \max\{m, S/p\}$ we have $\max\{m, S/p\} \leq B^* \leq 2 \cdot \max\{m, S/p\}$ which proves the approximation factor of 2.

Denote $w_{p+1} = x$. It remains to prove that the run of PROBE always succeeds, i.e., that the optimal bottleneck value of the sequence $w_1, w_2, \ldots w_p, w_{p+1}$ is at most $B = 2 \cdot \max\{m, S/p\}$ in every round. Indeed, if PROBE(B) does not succeed in creating p partitions, then $w_i + w_{i+1} > B \geq 2S/p$ must hold for all $1 \leq i \leq p$. But then:

$$S = \sum_{i=1}^{p+1} w_i \geq \sum_{i=1}^{\lfloor (p+1)/2 \rfloor} (w_{2i-1} + w_{2i}) > \lfloor (p+1)/2 \rfloor \cdot 2S/p \geq S,$$

a contradiction, which proves the correctness of the algorithm. The space requirement is dominated by the weights of the p partitions, yielding the bound $O(p \log(mn))$. □

Theorem 5. *There exists a 2-approximation algorithm for* PARTB *that uses* $O(\log(mn))$ *space.*

Proof. We simply compute in one pass the total weight S and the maximum m, then output $\max\{m, \frac{S}{p}\} + m$. By Lemma 1 we have $\max\{m, \frac{S}{p}\} \leq B^* \leq \max\{m, \frac{S}{p}\} + m$. Hence, the approximation ratio is at most $1 + m/\max\{m, \frac{S}{p}\} \leq 2$. The total weight of the stream is at most mn, therefore the space usage is $O(\log(mn))$. □

6 Space Lower Bounds

6.1 A Linear Space Lower Bound for Exact Algorithms

In this section, we show that any possibly randomized exact streaming algorithm for either PART or PARTB that performs one pass over the input requires $\Omega(n)$ space. We show this by a reduction from the INDEX problem in one-way two-party communication complexity.

Definition 1 (Index Problem). *Let $S = (S_1, \ldots, S_N)$ where $S \in \{0,1\}^N$, and let $I \in \{1, \ldots, N\}$. Alice is given S, Bob is given I. Alice sends message M to Bob and upon reception Bob outputs S_I.*

We consider a version of INDEX where the index I is chosen from the set $\{\lceil N/2 \rceil, \ldots, N\}$ uniformly at random. It is well-known [17] that the one-way randomized communication complexity of INDEX is $\Omega(N)$, and the modification in the input distribution restricting the index I to be chosen from the set $\{\lceil N/2 \rceil, \ldots, N\}$ does not change its hardness.

Lemma 7 (Hardness of the Index Problem). *If S is chosen uniformly at random from $\{0,1\}^N$, and I is chosen uniformly at random from the set $\{\lceil N/2 \rceil, \ldots, N\}$ and the failure probability of the protocol is at most $1/3$, then $\mathrm{Exp}_S|M| = \Omega(N)$.* □

Reduction. Given a streaming algorithm ALG that solves PART or PARTB on a stream of length at most $3n$ using space s, we specify a protocol for an arbitrary instance (S, I) of the one-way two-party communication problem INDEX with $|S| = n$, such that the message size is at most s.

Remember that Alice holds $S \in \{0,1\}^N$ and Bob holds $I \geq \lceil N/2 \rceil$. Our protocol is the following: Alice generates the sequence $Y \in \{1,3\}^{2N}$ such that $Y_i = 2 \cdot S_{i/2} + 1$ for even i, and $Y_i = 4 - Y_{i+1}$ for odd i. Bob generates the sequence $Z = \underbrace{4 \ldots 4}_{2I-N-1}\, 2$.

Alice runs ALG on the sequence Y with the number of partitions $p = 2$. Once Y is entirely processed, she sends the resulting memory state of ALG to Bob. Bob continues running ALG on Alice's final memory state and feeds the sequence Z into ALG. Observe that from the point of view of ALG it is as if the input stream were the concatenation of Y and Z. The message size of the protocol equals the space usage of ALG after processing Y. If ALG is an algorithm for PART then ALG outputs the separator s_1 that separates the two partitions. If s_1 is even then Bob outputs 0, and if s_1 is odd then Bob outputs 1. If ALG is an algorithm for PARTB then ALG outputs the optimal bottleneck value B. If $B = 4I - 1$ then Bob outputs 0, otherwise Bob outputs 1.

We prove that the above protocol is correct (i.e. the value returned by Bob is S_I), which immediately yields the space lower bound.

Theorem 6. *Any possibly randomized exact one-pass streaming algorithm for* PART *or* PARTB *requires space $\Omega(n)$.*

Proof. First observe that $\sum_i Y_i + \sum_i Z_i = 4 \cdot N + (2I - N - 1) \cdot 4 + 2 = 8I - 2$. Let s_1^* denote the optimal split position. Suppose that a perfect balancing is achieved and the optimal bottleneck value is $4I - 1$. Since for all $i = 1, \ldots, N$ we have $Y_{2i-1} + Y_{2i} = 4$, this can only be achieved if s_1^* is even and $Y_{s_1^*} = 1$ which implies that $S_{s_1^*/2} = S_I = 0$. Suppose now that a perfect balancing cannot be achieved. This can only happen if s_1^* is odd and $Y_{s_1^*-1} = 3$ which implies that $S_{(s_1^*-1)/2} = S_I = 1$. Thus, the protocol is correct in both cases, and ALG can be used to solve INDEX. Lemma 7 gives hence a lower bound for the space requirements of ALG. □

6.2 $\Omega(\frac{1}{\epsilon} \log n)$ Space Lower Bound for Approximation Algorithms

In this section, we prove an $\Omega(\frac{1}{\epsilon} \log n)$ space lower bound for one-pass streaming algorithms for PART that compute a $(1+\epsilon)$-approximation. We prove this lower bound in the one-way two-party communication setting for instances of PART with $m = 1$ and $p = 2$. Alice is given a sequence $Y \in \{0,1\}^n$ and Bob is given a sequence $Z \in \{0,1\}^n$, and they have to split the sequence $X = Y \circ Z$ into two parts. Alice sends a message to Bob, and upon reception, Bob outputs the separator. We describe now the hard input distribution.

Let t be an integer that is to be determined later. Alice's input and Bob's input are independent from each other and they are constructed as follows:

Alice's input Y is a sequence of length n with $2(t-1)$ leading 1s, followed by an arbitrary sequence of length $n-3t+2$, with elements from $\{0, 11\}$ (11 is a pair of ones), where the number of 11s is exactly t. Denote by $\mathcal{Y}$ the set of all such sequences. Then Y is chosen uniformly at random from $\mathcal{Y}$. Clearly, the weight of Y is $4t-2$, and $|\mathcal{Y}| = \binom{n-3t+2}{t}$.

Bob's input Z is a sequence of length n with the first $4(i-1)$ elements 1, and the remaining elements 0, for some $i \in \{1, 2, \ldots, t\}$. Denote all such sequences as $\mathcal{Z}$. Then Z is chosen uniformly at random from $\mathcal{Z}$. Observe that the weight of Z varies from 0 to $4(t-1)$, and $|\mathcal{Z}| = t$.

Note that an optimal partitioning of any $Y \circ Z$ instance splits one of the 11s in the second part of Alice's input.

Example: Let $t = 2$ and $n = 10$ and $p = 2$. Suppose that Alice holds $Y = 11\,00110110$. Bob's possible inputs are $Z_1 = 0000000000$ and $Z_2 = 1111000000$ of weight 0 and 4. The optimal partitioning of $Y \circ Z_1$ is $11\,001 \,|\, 10110\,0 \ldots 0$ and of $Y \circ Z_2$ is $11\,001101 \,|\, 10\,11110 \ldots 0$.

We give a lower bound on the space requirement of any possibly randomized communication protocol that solves instances of $\mathcal{Y} \times \mathcal{Z}$ exactly.

Lemma 8. *Any randomized one-way two-party communication protocol with error at most $\delta > 0$ that solves* Part *on instances of $\mathcal{Y} \times \mathcal{Z}$ has communication complexity at least*

$$\log\left(\frac{\binom{n-3t+2}{t}}{8\binom{t}{4\delta t} n^{4\delta t}}\right).$$

Proof. Let P be a randomized protocol as in the statement of the lemma. Then by Yao's Lemma [18], there is a deterministic protocol Q with distributional error at most δ that has the same communication complexity. We prove a lower bound on the communication complexity of Q.

Denote by $M_1, \ldots, M_k$ the possible messages from Alice to Bob, and let $\mathcal{Y}_i \subseteq \mathcal{Y}$ denote the set of inputs that Alice maps to message M_i. Note that for a fixed input for Bob, the protocol Q outputs the same result for all inputs in $\mathcal{Y}_i$. We define:

$$p_i = \Pr_{Y \leftarrow \mathcal{Y}_i, Z \leftarrow \mathcal{Z}}[Q \text{ errs on } (Y, Z)].$$

Since the distributional error of the protocol is δ, or in other words $\Pr_{Y \leftarrow \mathcal{Y}, Z \leftarrow \mathcal{Z}}[Q \text{ errs on } (Y, Z)] \leq \delta$, we obtain $\frac{\sum_i p_i |\mathcal{Y}_i|}{|\mathcal{Y}|} \leq \delta$. Let $i \in \{1, \ldots, l\}$ be the indices for which $p_i \leq 2\delta$. Then by the Markov Inequality, $\sum_{i=1}^{l} |\mathcal{Y}_i| \geq \frac{1}{2}|\mathcal{Y}|$.

We bound $|\mathcal{Y}_i|$ from above for all $i \in \{1, \ldots, l\}$. First, note that for a particular input $Z \in \mathcal{Z}$, the output of Q on (Y, Z) is the same for all $Y \in \mathcal{Y}_i$. Denote by $\mathcal{Y}_i^j$ the subset of $\mathcal{Y}_i$ such that for each $Y^j \in \mathcal{Y}_i^j : \Pr_{Z \leftarrow \mathcal{Z}}[Q \text{ errs on } (Y^j, Z)] = \frac{j}{t}$, or in other words, there are j inputs of Bob such that the protocol fails on Y^j, and for the remaining $t-j$ inputs of Bob, the protocol succeeds. Consider the set $\mathcal{Y}_i^0$, i.e., for each $Y \in \mathcal{Y}_i^0$, the protocol succeeds on any input of Bob. This determines all positions of the pairs of 1s in Alice's input, and therefore, there is only a single such element and we obtain $|\mathcal{Y}_i^0| \leq 1$. Similarly, we obtain:

$$|\mathcal{Y}_i^j| \leq \binom{t}{j} n^j,$$

since the protocol errs on at most j inputs of Bob, therefore the position of $t-j$ pairs of 1s is fixed and only j pairs of 1s may differ (we allow them to have an arbitrary position in Y which is a very rough estimate).

We apply the Markov Inequality again: for at least half of the elements of $\mathcal{Y}_i$, the protocol errs with probability at most 4δ. Therefore:

$$\frac{1}{2}|\mathcal{Y}_i| \leq \sum_{j \leq 4\delta t} |\mathcal{Y}_i^j| \leq \sum_{j \leq 4\delta t} \binom{t}{j} n^j \leq 2\binom{t}{4\delta t} n^{4\delta t},$$

and thus $|\mathcal{Y}_i| \leq 4\binom{t}{4\delta t} n^{4\delta t}$. This implies that:

$$l \geq \frac{|\mathcal{Y}|}{8\binom{t}{4\delta t} n^{4\delta t}} = \frac{\binom{n-3t+2}{t}}{8\binom{t}{4\delta t} n^{4\delta t}}.$$

Since the protocol sends at least l different messages, the communication complexity of the protocol is at least $\log(l)$, which implies the result. □

We choose t small enough so that a solution to any instance of $\mathcal{Y} \times \mathcal{Z}$ that is a $(1+\epsilon)$-approximation actually solves the instance exactly. This idea leads to our main lower bound theorem:

Theorem 7. *Any randomized one-way two-party communication protocol with error at most $\delta > 0$ (δ sufficiently small) that computes a $(1+\epsilon)$-approximation ($\frac{1}{\epsilon} = O(n^{1-\gamma})$ for any $\gamma > 0$) to* Part *on instances of $\mathcal{Y} \times \mathcal{Z}$ has communication complexity at least $\Omega\left(\frac{1}{\epsilon}\log n\right)$.*

Proof. We choose t small enough that a solution to any instance of $\mathcal{Y} \times \mathcal{Z}$ that is a $(1+\epsilon)$-approximation actually solves the instance exactly. Remark again that the weight of Y is $4t-2$ and the weight of Z is $4(i-1)$. Since the total weight is even, there is always a partitioning with weight $2t-1+2(i-1)$. Therefore, any partitioning that does not achieve an optimal balancing has an approximation factor of at least $\frac{2t-1+2(i-1)+1}{2t-1+2(i-1)}$, and we wish to choose t such that this approximation factor is worse than a $(1+\epsilon)$ approximation. Therefore, we have to choose t small enough such that for any $i \in \{1, 2, \ldots, t\}$

$$\frac{1}{2t-1+2(i-1)} > \epsilon,$$

which implies that $t < \frac{1}{4\epsilon} + \frac{3}{4}$. We choose $t = \frac{1}{4\epsilon}$ and plug this value into the communication lower bound from Lemma 8. Using standard bounds on binomial coefficients:

$$\begin{aligned}
\Omega\left(\log\left(\frac{\binom{n-3t+2}{t}}{8\binom{t}{4\delta t} n^{4\delta t}}\right)\right) &= \Omega\left(\log\left(\frac{\left(4\epsilon(n-\frac{3}{4\epsilon}+2)\right)^{\frac{1}{4\epsilon}}}{8n^{\delta/\epsilon}\left(\frac{e}{4\delta}\right)^{\delta/\epsilon}}\right)\right) \\
&= \Omega\left(\frac{1}{4\epsilon}\log(4\epsilon n - 3 + 8\epsilon) - \frac{\delta}{\epsilon}\log(\frac{ne}{4\delta})\right) \\
&= \Omega\left(\frac{1}{4\epsilon}\log(4\epsilon n) - \frac{\delta}{\epsilon}\log(\frac{ne}{4\delta})\right) \\
&= \Omega(\left(\frac{1}{\epsilon}\log n\right),
\end{aligned}$$

for a sufficiently small but constant δ, and $\epsilon = \mathrm{O}(n^{1-\gamma})$ for any $\gamma > 0$. This proves the result. □

7 Conclusion and Open Problems

In this paper, we presented one-pass $(1+\epsilon)$-approximation streaming algorithms for partitioning integer sequences that are based on the parametric search framework. We designed a new method for carrying out feasibility tests of multiple parameters simultaneously, leading to an improvement over the naïve application of the method. We compromised our algorithms with lower bounds showing that an optimal solution cannot be computed with sublinear space, and a $(1+\epsilon)$-approximation requires space $\Omega(\frac{1}{\epsilon})$, rendering our algorithms almost tight with respect to the dependency on parameter ϵ.

We demonstrated that the parametric search framework can successfully be applied in the streaming setting, and even though the streaming model is very restrictive, it allows an improvement over the naïve application of the method. We believe that other problems admit parametric search algorithms in the streaming setting, and we leave the identification and the study of those as an open problem.

The most intriguing open question concerns the situation where a streaming algorithm has no information about the problem parameters m, n, and S, the maximal weight of an element of the stream, the stream length, and the total weight of the stream, respectively. For this situation, we designed a 2-approximation algorithm, however, there is no argument contradicting the existence of a $(1+\epsilon)$-approximation algorithm.

References

1. Bokhari, S.H.: Partitioning problems in parallel, pipeline, and distributed computing. IEEE Trans. Comput. **37**(1) (1988) 48–57
2. Hansen, P., Lih, K.W.: Improved algorithms for partitioning problems in parallel, pipelined, and distributed computing. IEEE Trans. Comput. **41**(6) (1992) 769–771
3. Manne, F., Olstad, B.: Efficient partitioning of sequences. IEEE Trans. Comput. **44**(11) (1995) 1322–1326
4. Manne, F., Sørevik, T.: Optimal partitioning of sequences. J. Algorithms **19**(2) (1995) 235–249
5. Khanna, S., Muthukrishnan, S., Skiena, S.: Efficient array partitioning. In: Automata, Languages and Programming. Volume 1256., Springer Berlin Heidelberg (1997) 616–626
6. Han, Y., Narahari, B., Choi, H.A.: Mapping a chain task to chained processors. Inf. Process. Lett. **44**(3) (1992) 141–148
7. Miguet, S., Pierson, J.M.: Heuristics for 1d rectilinear partitioning as a low cost and high quality answer to dynamic load balancing. In: Proceedings of the International Conference and Exhibition on High-Performance Computing and Networking. HPCN Europe '97, London, UK, UK, Springer-Verlag (1997) 550–564
8. Pinar, A., Aykanat, C.: Fast optimal load balancing algorithms for 1d partitioning. J. Parallel Distrib. Comput. **64**(8) (2004) 974–996
9. Frederickson, G.N.: Optimal algorithms for tree partitioning. In: Proceedings of the Second Annual ACM-SIAM Symposium on Discrete Algorithms. SODA '91, Philadelphia, PA, USA, Society for Industrial and Applied Mathematics (1991) 168–177
10. Schamberger, S., Wierum, J.M.: Partitioning finite element meshes using space-filling curves. Future Gener. Comput. Syst. **21**(5) (2005) 759–766
11. Konrad, C.: Two-constraint domain decomposition with space filling curves. Parallel Comput. **37**(4-5) (2011) 203–216
12. Bader, M.: Space-Filling Curves - An Introduction with Applications in Scientific Computing. Volume 9 of Texts in Computational Science and Engineering. Springer-Verlag (2013)
13. Muthukrishnan, S.: Data streams: Algorithms and applications. In: Foundations and Trends in Theoretical Computer Science. Now Publishers Inc (2005)
14. Iqbal, M.A.: Approximate algorithms for partitioning problems. International Journal of Parallel Programming **20**(5) (1991) 341–361
15. Megiddo, N.: Combinatorial optimization with rational objective functions. In: Proceedings of the Tenth Annual ACM Symposium on Theory of Computing. STOC '78, New York, NY, USA, ACM (1978) 1–12
16. Megiddo, N.: Applying parallel computation algorithms in the design of serial algorithms. J. ACM **30**(4) (1983) 852–865
17. Kushilevitz, E., Nisan, N.: Communication complexity. Cambridge University Press (1997)

18. Yao, A.C.C.: Probabilistic computations: Toward a unified measure of complexity. In: Proceedings of the 18th Annual Symposium on Foundations of Computer Science. FOCS '77, Washington, DC, USA, IEEE Computer Society (1977) 222–227